arXiv:1412.2470v1 [cs.CC] 8 Dec 2014

Bounded Treewidth and Space-efficient Linear Algebra

Nikhil Balaji & Samir Datta*

Chennai Mathematical Institute (CMI), India
{nikhil,sdatta}@cmi.ac.in

Abstract. Motivated by a recent result of Elberfeld, Jakoby and Tantau[EJT10] showing that MSO properties are Logspace computable on graphs of bounded tree-width, we consider the complexity of computing the determinant of the adjacency matrix of a bounded tree-width graph and as our main result prove that it is in Logspace. It is important to notice that the determinant is neither an MSO-property nor counts the number of solutions of an MSO-predicate. This technique yields Logspace algorithms for counting the number of spanning arborescences and directed Euler tours in bounded tree-width digraphs.
We demonstrate some linear algebraic applications of the determinant algorithm by describing Logspace procedures for the characteristic polynomial, the powers of a weighted bounded tree-width graph and feasibility of a system of linear equations where the underlying bipartite graph has bounded tree-width.
Finally, we complement our upper bounds by proving L-hardness of the problems of computing the determinant, and of powering a bounded tree-width matrix. We also show the GapL-hardness of Iterated Matrix Multiplication where each matrix has bounded tree-width.

1 Introduction

The determinant is a fundamental algebraic invariant of a matrix. For an $n \times n$ matrix A the determinant is given by the expression $\mathrm{Det}(A) = \sum_{\sigma \in S_n} \mathrm{sign}(\sigma) \prod_{i \in [n]} a_{i,\sigma(i)}$ where S_n is the symmetric group on n elements, σ is a permutation from S_n and $\mathrm{sign}(\sigma)$ is the parity of the number of inversions in σ ($\mathrm{sign}(\sigma) = 1$ if the number of inversions in σ is even and 0 if it is odd). Even though the summation in the definition runs over $n!$ many terms, there are many efficient sequential [vzGG13] and parallel [Ber84] algorithms for computing the determinant.

Apart from the inherently algebraic methods to compute the determinant there are also combinatorial algorithms (see, for instance, Mahajan and Vinay [MV97]) which extend the definition of determinant as a signed sum of cycle covers in the weighted adjacency matrix of a graph. [MV97] are thus able to give another proof of the GapL-completeness of the determinant, a result first proved

* Part of the work was done on a visit to the Institute for Theoretical Computer Science at Leibniz University Hannover

by Toda [Tod91]. For a more complete discussion on the known algorithms for the determinant, see [MV97].

Armed with this combinatorial interpretation of the determinant and faced with its GapL-hardness, one can ask if the determinant is any easier when the underlying matrix represents simpler classes of graphs. Datta, Kulkarni, Limaye, Mahajan [DKLM10] study the complexity of the determinant and permanent, when the underlying directed graph is planar and show that they are as hard as the general case - GapL and #P-hard, respectively. We revisit these questions in the context of bounded tree-width graphs.

Many NP-complete graph problems become tractable when restricted to graphs of bounded tree-width. In an influential paper, Courcelle [Cou90] proved that any property of graphs expressible in Monadic Second Order MSO logic can be decided in linear time on bounded tree-width graphs. For example, Hamiltonicity is an MSO property and hence deciding if a bounded tree-width graph has a Hamiltonian cycle can be done in linear time. More recently Elberfeld, Jakoby, Tantau [EJT10] showed that in fact, MSO properties on bounded tree-width graphs can be decided in L.

We study the Determinant problem when the underlying directed graph has bounded tree-width and show a Log-space upper bound. In the same vein we also compute other linear algebraic invariants of a bounded tree-width matrix, such as the characteristic polynomial, rank and powers of a matrix in Logspace. Interpreting rectangular matrices as (weighted) bipartite graphs, we are also able to show that checking for the feasibility of a system of linear equations for such matrices arising from bounded tree-width bipartite graphs is in L. FSLE has previously been studied for general graphs in [ABO99] where it is shown to be complete for the first level of the Logspace counting hierarchy: $\mathsf{L}^{\mathsf{C}=\mathsf{L}}$.

We give a tight bound on the complexity of the determinant by showing that it is L-hard via a reduction from directed reachability in paths. We also show that it is unreasonable to attempt to extend the Logspace upper bound of determinant and powering to Iterated Matrix Multiplication (IMM) of bounded tree-width matrices, by showing GapL-hardness for IMM. It is worthwhile to contrast this with the case of general graphs, where the Determinant, IMM and Matrix Powering are known to be inter-reducible to each other and hence complete for GapL.

Counting spanning trees in directed graphs is easily seen to be in GapL by the matrix tree theorem[Sta13]. Counting modulo 2 has recently been proved to be $\oplus$L-hard for planar graphs in [DK14]. As a direct consequence of our determinant result and the Kirchoff matrix tree theorem it follows that the problem is in L for graphs of bounded tree-width.

The **BEST** theorem due to De Bruijn, Ehrenfest, Smith and Tutte gives an exact formula for the number of Euler tours in a directed graph (see Fact 3) in terms of the number of directed spanning trees of the graph.

1.1 Our Results and Techniques

Through out this paper, we work with matrices with entries from $\mathbb{Q}$, unless stated otherwise. We show that the following can be computed/tested in L:

1. (Main Result) The Determinant of an $(n \times n)$ matrix A whose underlying undirected graph has bounded tree-width. As a corollary we can also compute the coefficients of the characteristic polynomial of a matrix.
2. The inverse of an $(n \times n)$ matrix A whose underlying undirected graph has bounded tree-width. As a corollary we get a Logspace algorithm to compute the powers A^k of a matrix A (with rational entries) whose support is a bounded tree-width digraph.
3. Testing if a system of rational linear equations $Ax = b$ is feasible where A is (a not necessarily square) matrix whose support is the biadjacency matrix of an undirected bipartite graph of bounded tree-width.
4. The number of Spanning Trees in graphs of bounded tree-width.
5. The number of Euler tours in a bounded tree-width directed graph

We also show hardness results to complement the above easiness results:

1. Computing the determinant of a bounded tree-width matrix is L-hard which precludes further improvement in the Logspace upper bound.
2. Computing the iterated matrix multiplication of bounded tree-width matrices is GapL-hard which precludes attempts to extend the L-bound on powering matrices of bounded tree-width to iterated matrix multiplication.
3. Powering matrices are however L-hard which prevents attempts to further improve the L-bound on matrix powering.

At the core of the results is our algorithm to compute the determinant by writing down an MSO formula that evaluates to true on every valid cycle cover of the bounded tree-width graph underlying A. The crucial point being that the cycle covers are parameterised on the number of cycles in the cycle cover, a quantity closely related to the sign of the cycle covers. This makes it possible to invoke the cardinality version of Courcelle's theorem(for Logspace) due to [EJT10] to compute the determinant. A more subtle point is that in order to keep track of the number of cycles as the size of a set of vertices, we need to pick one vertex per cycle. Picking one vertex per cycle is done by choosing the "smallest" vertex in the cycle. In order to pick a vertex in a cycle cover, we need to define a total order on the vertices which makes this part of the proof technically challenging.

We use this determinant algorithm and the Kirchoff matrix tree theorem along with the BEST theorem to count directed Euler tours.

1.2 Organization of the paper

Section 2 introduces some notation and terminology required for the rest of the paper. In Section 3, we give a Logspace algorithm to compute the Determinant

of matrices of bounded tree-width and give some linear algebraic and graph theoretic applications. In Section 4, we give some L-hardness results to complement our Logspace upperbounds. In Section 5, we mention some problems that remain open.

2 Preliminaries

2.1 Background on Graph Theory

Definition 1. *Given an undirected graph $G = (V_G, E_G)$ a tree decomposition of G is a tree $T = (V_T, E_T)$ (the vertices in $V_T \subseteq 2^{V_G}$ are called bags), such that*

1. *Every vertex $v \in V_G$ is present in at least one bag, i.e., $\cup_{X \in V_T} X = V_G$.*
2. *If $v \in V_G$ is present in bags $X_i, X_j \in V_T$, then v is present in every bag X_k in the unique path between X_i and X_j in the tree T.*
3. *For every edge $(u, v) \in E_G$, there is a bag $X_r \in V_T$ such that $u, v \in X_r$.*

The width of a tree decomposition is the $\max_{X \in V_T}(|X| - 1)$. The tree width of a graph is the minimum width over all possible tree decomposition of the graph.

Definition 2. *Given a weighted directed graph $G = (V, E)$ by its adjacency matrix $[a_{ij}]_{i,j \in [n]}$, a cycle cover $\mathcal{C}$ of G is a set of vertex-disjoint cycles that cover the vertices of G. I.e., $\mathcal{C} = \{C_1, C_2, \ldots, C_k\}$, where $V(C_i) = \{c_{i_1}, \ldots, c_{i_r}\} \subseteq V$ such that (c_{i_1}, c_{i_2}), (c_{i_2}, c_{i_3}), …, $(c_{i_{r-1}}, c_{i_r})$, $(c_{i_r}, c_{i_1}) \in E(C_i) \subseteq E$ and $\sqcup_{i=1}^{k} V(C_i) = V$.*

Fact 1 *The weight of the cycle $C_i = \prod_{j \in [r]} wt(a_{ij})$ and the weight of the cycle cover $wt(\mathcal{C}) = \prod_{i \in [k]} wt(C_i)$. The sign of the cycle cover $\mathcal{C}$ is $(-1)^{n+k}$.*

Every permutation $\sigma \in S_n$ can be written as a union of vertex disjoint cycles. Hence a permutation corresponds to a cycle cover of a graph on n vertices. In this light, the determinant of an $(n \times n)$ matrix A can be seen as a signed sum of cycle covers:

$$\det(A) = \sum_{\text{cycle cover } \mathcal{C}} \text{sign}(\mathcal{C}) \text{wt}(\mathcal{C})$$

2.2 Background on MSO-logic

Definition 3 (Monadic Second Order Logic). *Let the variables $V = \{v_1, v_2, \ldots, v_n\}$ denote the vertices of a graph $G = (V, E)$. Let X, Y denote subsets[1] of V or E. Let $E(x, y)$ be the predicate that evaluates to 1 when there is an edge between x and y in G. A logical formula ϕ is called an MSO-formula if it can be constructed using the following:*

[1] The case when quantification over subset of edges is not allowed is referred to as MSO_1 which is known to be strictly less powerful than MSO_2, the case when edge set quantification is allowed. Throughtout our paper, we will work with MSO_2 and hence we will just refer to it as MSO.

- $v \in X$
- $v_1 = v_2$
- $E(v_1, v_2)$
- $\phi_1 \vee \phi_2$, $\phi_1 \wedge \phi_2$, $\neg\phi$
- $\exists x\phi, \forall x\phi$
- $\exists X\phi, \forall X\phi$

In addition, if the Gaifman graph[2] *of the relation is bounded treewidth then we can use any predicate in item 3 above. A property Π of graphs is* MSO*-definable, if it can be expressed as a* MSO *formula ϕ such that ϕ evaluates to TRUE on a graph G if and only if G has property Π.*

Definition 4 (Solution Histogram). *Given a graph $G = (V, E)$ and an* MSO *formula $\phi(X_1, \ldots, X_d)$ in free variables $X_1, \ldots, X_d$, where $X_i \subseteq V$ (or E), the $(i_1, \ldots, i_d)$-th entry of histogram(G, ϕ) gives the number of subsets $S_1, \ldots, S_d$ such that $|S_j| = i_j$ for which $\phi(S_1, \ldots, S_d)$ is true.*

We need the following results from [EJT10]:

Theorem 1 (Logspace version of Bodlaender's theorem). *For every $k \geq 1$, there is a Logspace machine that on input of any graph G of tree width at most k outputs a width-k tree decomposition of G.*

Theorem 2 (Logspace version of Courcelle's theorem). *For every $k \geq 1$ and every* MSO*-formula ϕ, there is a Logspace machine that on input of any logical structure $\mathcal{A}$ of tree width at most k decides whether $A \models \phi$ holds.*

Theorem 3 (Cardinality version of Courcelle's theorem). *Let $k \geq 1$ and let $\phi(X_1, \ldots, X_d)$ be an* MSO*-formula on free variables $X_1, \ldots, X_d$. Then there is a Logspace machine that on input of the tree decomposition of a graph G of treewidth at most k,* MSO*-formula ϕ and $(i_1, \ldots, i_d)$, outputs the value of histogram(G, ϕ) at $|X_1| = i_1, \ldots, |X_d| = i_d$.*

3 Determinant Computation

Given a square $\{0, 1\}$-matrix A, we can view it as the bipartite adjacency matrix of a bipartite graph H_A. The permanent of this matrix A counts the number of perfect matchings in H_A, while the determinant counts the signed sum of perfect matchings in H_A.

If G is a bounded treewidth graph then we can count the number of perfect matchings in G in L[EJT10] (see also [DDN13]). Hence the complexity of the permanent of A, above is well understood in this case while the complexity of computing the determinant is not clear.

[2] The Gaifman graph (also called the *Primal Graph*) of a binary relation $R \subseteq A \times A$ is the graph whose nodes are elements of A and an edge joins a pair of variables x, y if $(x, y) \in R$.

On the other hand the determinant of a $\{0,1\}$-matrix reduces (say by a reduction g_{MV}) to counting the number of paths in another graph (see e.g. [MV97]). Also counting s,t-paths in a bounded treewidth graph is again in L via [EJT10] (see also [DDN13]). But the problem with this approach is that that the graph $g_{MV}(G)$ obtained by reducing a bounded treewidth G is not bounded treewidth.

However, we can also view A as the adjacency matrix of a directed graph G_A. If G_A has bounded treewidth (which implies that H_A also has bounded treewidth, see Proposition 4 in Appendix B) then we have a way of computing the determinant of A. To see this, consider the following lemma:

Lemma 1. *There is an* MSO*-formula* $\phi(X,Y)$ *with free variables* X,Y *that take values from the set of subsets of vertices and edges respectively, such that* $\phi(X,Y)$ *is true exactly when* X *is the set of heads of a cycle cover* Y *of the given graph.*

Before proving this Lemma we need some preprocessing. Let G be the input graph of bounded tree-width. We will augment G with some new vertices and edges to yield a graph G' again with a tree decomposition T' of bounded tree-width. Then we have:

Lemma 2. *There exists a relation* NXT *on vertices of* G' *which satisfies the following:*

1. NXT *is compatible*[3] *with the tree decomposition* T'
2. NXT *is a partial order on the vertices of* G'
3. NXT *is computable in* L
4. *The transitive closure* NXT^* *is a total order when restricted to the vertices of* G
5. NXT^* *is expressible as an* MSO*-formula over the vocabulary of* G' *along with* NXT.

The construction of such a relation is fairly straight forward and considered folklore in the Finite Model Theory literature (See for example Proposition VI.4 in [CF12]. Here we include an proof of Lemma 2 (obtained independently) in the appendix for the sake of completeness.

Proof. (of Lemma 1) We write an MSO formula ϕ on free variables X,Y, such that $Y \subseteq E$ and $X \subseteq V$, such that ϕ evaluates to true on any set of heads of a cycle cover S. The MSO predicate essentially verifies that the subgraph induced by Y indeed forms a cycle cover of G. Our MSO formula is of the form [4]:

$$\phi(X,Y) \equiv (\forall v \in V)(\exists! h \in X)[\mathsf{DEG}(v,Y) \wedge \mathsf{PATH}(h,v,Y) \wedge (\mathsf{NXT}^*(h,v) \vee \mathsf{EQ}(h,v))]$$

where,

[3] Binary relation R is said to be compatible with the tree decomposition T' of G if the Gaifman graph of R has T' as its tree decomposition.

[4] Note that since we require that for a given X,Y, every $v \in V$ has a unique $h \in X$, our formula is not monotone, i.e., If $X \subseteq X'$ are two sets of heads then if $\phi(X,Y)$ is true doesn't imply $\phi(X',Y)$ is also true (consider vertices in $X' \setminus X$, since $X' \subseteq X$, they will have two different h, h' such that the PATH and NXT^* predicates are true contradicting uniqueness of h

1. $\mathsf{DEG}(v, Y)$ is the predicate that says that the in-degree and out-degree of v (in the subgraph induced by the edges in Y) is 1.
2. $\mathsf{PATH}(x, y, Y)$ is the predicate that says that there is a path from x to y in the graph induced by edges of Y.
3. $\mathsf{EQ}(h, v) = 1$ iff $h = v$

One can check that all the predicates above are MSO-definable.

Lemma 1 along with the Fact 1 yields:

Lemma 3. *Given an $(n \times n)$ bounded treewidth matrix A with integer entries, there is a Logspace algorithm that constructs an $(m \times m)$(where $m = \mathsf{poly}(n)$) matrix B with entries from $\{0, 1\}$, such that $det(A) = det(B)$ and the treewidth of B is the same as the treewidth of A.*

Thus, using the histogram version of Courcelle's theorem from [EJT10] and Lemma 3, we get:

Theorem 4. *The determinant of a matrix A with integer entries, which can be viewed as the adjacency matrix of a weighted directed graph of bounded treewidth, is in* L.

Proof. Firstly, obtain the matrix B from A using Lemma 3. The histogram version of Courcelle's theorem as described in [EJT10] when applied to the formula $\phi(X, Y)$ above yields the number of cycle covers of G_B parametrized on $|X|, |Y|$. But in the notation of Fact 1 above, $|X| = k$ and $|Y| = n$, so we can easily compute the determinant as the alternating sum of these counts.

Corollary 1. *There is a Logspace algorithm that takes as input a $(n \times n)$ bounded treewidth matrix A, 1^m, where $1 \leq m \leq n$ and computes the coefficient of x^m in the characteristic polynomial ($\chi_A(x) = det(xI - A)$) of A.*

The characteristic polynomial of an $(n \times n)$ matrix A is the determinant of the matrix $A(x) = xI - A$. We could use Theorem 4 to compute this quantity (since $A(x)$ is bounded treewidth, if A is bounded treewidth). However, Theorem 4 holds only for matrices with integer entries while the matrix $A(x)$ contains entries in the diagonal involving the indeterminate x.

We proceed as follows: In the directed graph corresponding to A, replace a self loop on a vertex of weight $x - d$ by a gadget of weight $-d$ in parallel with a self loop of weight x (In the event that there is no self loop on a vertex in A, add a self loop of weight x on the vertex). Replace the weights on the other edges according to the gadget in Lemma 3. We have added exactly n self loops, each of weight x (for the original vertices of A).

We first consider a generalisation of the determinant of $\{0, 1\}$-matrices of bounded tree-width viz. the determinant of matrices where the entries are from a set whose size is a fixed universal constant and the underlying graph consisting of the non-zero entries of A is of bounded tree-width.

Lemma 4. *Let A be a matrix whose entries belong to a set S of fixed size independent of the input or its length. If the underlying digraph with adjacency matrix A', where $A'_{ij} = 1$ iff $A_{ij} \neq 0$, is of bounded tree-width then the determinant of A can be computed in* L.

Proof. Let $s = |S|$ be a universal constant, $S = \{c_1, \ldots, c_s\}$ and let val_i be the predicates that partitions the edges of G according to their values i.e. $\mathrm{val}_i(e)$ is true iff the edge e has value $c_i \in S$. Our modified formula $\psi(X, Y_1, \ldots, Y_s)$ will contain s unquantified new edge-set variables $Y_1, \ldots, Y_s$ along with the old vertex variable X, and is given by:

$$\forall e \in E((e \in Y_i \rightarrow \mathrm{val}_i(e)) \wedge (e \in Y \leftrightarrow \vee_{i=1}^{s}(e \in Y_i) \wedge \phi(X, Y))$$

Notice that we verify that the edges in the set Y_i belong to the i^{th} partition and each eadge in Y is in one of the Y_i's. The fact that the Y_i's form a partition of Y follows from the assumption that $\mathrm{val}_i(e)$ is true for exactly one $i \in [s]$ for any edge e.

To obtain the determinant we consider the histogram parameterised on the s variables $Y_1, \ldots, Y_s$ and the heads X. For an entry indexed by $x, y_1, \ldots, y_s$, we multiply the entry by $(-1)^{n+x} \prod_{i=1}^{s} {c_i}^{y_i}$ and take a sum over all entries.

In light of the Lemma above, we can compute the characteristic polynomial as follows:

Proof. (of Corollary 1) While counting the number of cycle covers with k cycles, we can keep track of the number of self-loops occurring in a cycle cover. It is easy to see that we can obtain the coefficient of x^r in the characteristic polynomial from the histogram outlined in Lemma 4. Hence we can also compute the characteristic polynomial in L.

Corollary 2. *There is a Logspace algorithm that takes as input a bounded treewidth $(n \times n)$ matrix A, and computes the rank of A.*

Proof. Compute the characteristic polynomial of A and use the fact that the rank of A is a number r such that x^{n-r} is the smallest power of x wth a non-zero coefficient.

$\mathsf{FSLE}(A, b)$ is the following problem: Given a system of m linear equations (with integer coefficients, w.l.o.g) in variables $z_1, \ldots, z_n$ and a target vector b, we want to check if there is a feasible solution to $Az = b$. That is, we want to decide if there is a setting of the variable vector $z \in \mathbb{Q}^n$ such that, $Az = b$ holds for a bounded treewidth matrix $A \in \mathbb{Z}^{m \times n}$ (when we say a rectangular matrix is bounded treewidth, we mean the underlying bipartite graph on $(m+n)$ vertices has bounded treewidth).

Corollary 3. *For a bounded treewidth matrix $A_{m \times n}$ and vector $b_{n \times 1}$, $\mathsf{FSLE}(A, b)$ is in* L.

Proof. A can be interpreted as the biadjacency matrix of a bipartite graph. Now, consider the matrix $B = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$ – this is a matrix of dimension $(m+n) \times (m+n)$. It is easy to see that B corresponds to the adjacency matrix of A. Let row-rank(A) = column-rank$(A) = r$. Since A and A^T have the same rank, $\text{rank}(A) + \text{rank}(A^T) = 2r = \text{rank}(B)$. Therefore, in order to find the rank of the rectangular matrix A, we can use the Logspace procedure for matrix rank given by Corollary 2. Now, we know that the system of linear equations given by A, b is feasible if and only if $\text{rank}(A) = \text{rank}([A : b])$.

Corollary 4. *There is a Logspace algorithm that takes as input a $(n \times n)$ bounded treewidth matrix A, $1^i, 1^j, 1^k$ and computes the k-th bit of A^{-1}_{ij}.*

Proof. The inverse of a matrix A is the matrix $B = \frac{\mathbf{C}^T}{det(A)}$ where $\mathbf{C} = (C_{ij})_{1 \le i,j \le n}$ is the cofactor matrix, whose (i,j)-th entry $C_{ij} = (-1)^{i+j}\det(A_{ij})$ is the determinant of the $(n-1) \times (n-1)$ matrix obtained from A by deleting the i-th row and j-th column. If we can compute C_{ij} in L, we can compute the entries of B via integer division which is known to be in L from [HAB02]. To this end, consider the directed graph G_A represented by A. To compute $\det(A_{ij})$, swap the columns of A such that the j-th column becomes the i-th column. The graph so obtained is of bounded treewidth (To see this, notice that the swapping operation just re-routes all incoming edges of j to i and those of i to j. The tree decomposition of this graph is just obtained by adding vertices (i,j) to every bag in the tree decomposition of G_A and also the edges rerouted to the respective bags. This increases the treewidth by 2). Now, remove the i-th vertex in G_A and all edges incident to it to get a graph $G_{A'_{ij}}$ on $(n-1)$ vertices. The swapping operation changes the determinant of A_{ij} by a sign that is $(-1)^{i-j} = (-1)^{i+j}$. Computing the determinant of this modified matrix A'_{ij} yields C_{ij} as required. Since A'_{ij} is obtained from A by removing a vertex and all the edges incident on it, the treewidth of A'_{ij} is at most the treewidth of A. By Theorem 4, C_{ij} is in FL.

Corollary 5. *There is a Logspace algorithm that on input an $(n \times n)$ bounded treewidth matrix A, $1^m, 1^i, 1^j, 1^k$ gives the k-th bit of (i,j)-th entry of A^m.*

Proof. Consider $A' = (I - tA)^{-1}$ where I is the $(n \times n)$ identity matrix and t is a small constant to be chosen later. Notice that $A' = (I - tA)^{-1} = \sum_{j \ge 0} t^j A^j$. By choosing t as a suitably small power of 2 (say $2^{-p} = t$ such that $2^p > \|A\|$) and computing A' to a suitable accuracy, we can read the (i,j)-th entry of A^m off the appropriate bit positions of the (i,j)-th entry of A'. So, in essence the problem of powering bounded treewidth matrix A reduces to the problem of computing the inverse of a related matrix which is known to be in L via Corollary 4.

3.1 Spanning Trees and Directed Euler Tours

Fact 2 *The number of arborescences of a digraph equals any cofactor of its Laplacian.*

where the Laplacian of a directed graph G is $D - A$ where D is the diagonal matrix with the D_{ii} being the out-degree of vertex i and A is the adjacency matrix of the underlying undirected graph. The BEST Theorem states:

Fact 3 ([AEB87][TS41]) *The number of Euler Tours in a directed Eulerian graph K is exactly:*

$$t(K) \prod_{v \in V} (deg(v) - 1)!$$

where $t(K)$ is the number of arborescences in K rooted at an arbitrary vertex of K and $deg(v)$ is the indegree as well as the outdegree of the vertex v.

We combine Facts 2 and 3with Theorem 4 to compute the number of directed Euler Tours in a directed Eulerian graph in L.

Use the Kirchoff Matrix Tree theorem[Sta13] and Fact 3:

Corollary 6. *Counting arborescences and directed Euler Tours in a directed Eulerian graph G (where the underlying undirected graph is bounded treewidth) is in L.*

4 Hardness Results

We show a couple of hardness results to complement our Logspace upper bounds.

Proposition 1 (Hardness of Bounded Treewidth Determinant). *For all constant $k \geq 1$, computing the determinant of an $(n \times n)$ matrix A whose underlying undirected graph has treewidth at most k is L-hard.*

Proof. We reduce the problem ORD of deciding for a directed path P and two vertices $s, t \in V(P)$ if there is a path from s to t (known to be L-complete via [Ete97]) to computing the determinant of bounded treewidth matrices (Note that P is a path and hence it has treewidth 1). Our reduction is as follows: Given a directed path P with source a, sink b and distinguished vertices s and t, we construct a new graph P' as follows: Add edges $(a, s'), (s', t), (t, s), (s, a)$ and (b, t') and remove edges $(s', s), (t, t')$ where s' and t' are vertices in P such that $(s', s), (t, t') \in E(P)$ (See Figure 1).

We claim that there is a directed path between s and t if and only if the determinant of the adjacency matrix of P' is zero. If there is a directed path from s to t in P, then there are two cycle covers in $P' : (a, s')(s, t)(t', b)$, with three cycles and $(a, s', t, s), (t', b)$, with two cycles. Using Fact 1, the signed sum of these cycle covers is $(-1)^{n+3} + (-1)^{n+2} = 0$, which is the determinant of P'.

In the case that P has a directed path from t to s (see Figure 2), then there is one cycle namely (a, t, s', b, t', s). We argue as follows: The edges $(t, s), (s, b), (b, t'), (t', s')$ are in the cycle cover since they are the only incoming edges to s, b, t', s' respectively. So (t, s, b, t', s') is a part of any cycle cover of the graph. This forces one to pick the edge (s', a) and hence we have one cycle in the cycle cover for P'.

Fig. 1. s occurs before t

Fig. 2. t occurs before s

Proposition 2 (Hardness of Bounded Treewidth Matrix Powering). *Matrix Powering is* L*-hard under disjunctive truth table reductions.*

Proof. (of Proposition 2) We reduce ORD to matrix powering. Given an directed path P on n vertices and distinguished vertices s and t, we argue as follows: There is a directed path between s and t, then it must be of length i for an unique $i \in [n]$. Consider the matrix $(I + A_P)^n$: s and t are connected by a path if and only if $(I + A_P)^n_{s,t} \neq 0$. This is because $(I + A_P)^n_{s,t}$ gives the walks from s to t, and if at all there is a path from s to t, then there is definitely a walk of length at most n between them. Checking if this entry is zero can be done by a DNF which takes as input the bits of $(I + A_P)^n_{s,t}$.

Proposition 3 (Hardness of Bounded Treewidth IMM). *Given a sequence of bounded treewidth matrices with rational entries* $M_1, M_2, \ldots, M_n$ *and* $1^i, 1^j, 1^k$ *as input, computing the* k*-th bit of* (i, j)*-th entry of* $\prod_{l=1}^{n} M_l$ *is* GapL*-hard.*

Proof. (of Proposition 3) We reduce integer matrix powering to iterated matrix multiplication of bounded treewidth graphs. Given an $(n \times n)$ matrix A with n-bit entries, and $1^m, 1^i, 1^j$ the matrix powering problem is to find the (i, j)-th entry of A^m. From the underlying digraph $G_A = (V = \{v_1, v_2, \ldots, v_n\}, E)$, we construct a sequence of bounded treewidth matrices as follows: We construct two gadgets $\mathcal{V}_l$ and $\mathcal{V}_u$ – Both are graphs on $2n^2$ vertices divided in to n^2 partitions where each partition is a copy of V (such that there are no edges between vertices in the partition): $U = \sqcup_{i=1}^{n} U_i$ and $L = \sqcup_{i=1}^{n} L_i$ where each $L_i = U_i = V$. We also have the edges between:

1. $v_i \in U_j$ and $v_i \in U_{j+1}$
2. $v_i \in L_j$ and $v_i \in L_{j+1}$

for all $i \in [n], j \in [n-1]$.The edges are basically an identity perfect matching between U_i and U_{i+1} and also L_i and L_{i+1}. Now we add edges in $\mathcal{V}_l$ and $\mathcal{V}_u$ according to edges present in G_A: If $(v_i, v_1), \ldots, (u, v_r)$ are edges in G_A out of vertex v_i, then

1. In $\mathcal{V}_l$, we add an edge between $v_i \in L_i$ to $v_1, \ldots, v_r \in U_{i+1}$.
2. In $\mathcal{V}_u$, we add an edge between $v_i \in U_i$ to $v_{i_1}, \ldots, v_{i_r} \in L_{i+1}$

We now construct a walk gadget $\mathcal{W}$ using alternating copies of the vertex gadgets $\mathcal{V}_l$ and $\mathcal{V}_u$. To raise A to the m-th power, construct: $\mathcal{V}_1, \ldots, \mathcal{V}_m$ where $\mathcal{V}_1 = \mathcal{V}_l$, $\mathcal{V}_2 = \mathcal{V}_u$ and so on. Connect $\mathcal{V}_i$ and $\mathcal{V}_{i+1}$ by the following edges:

1. $v_i \in U_n$ and $v_i \in U_1, \forall i \in [n]$.
2. $v_i \in L_n$ and $v_i \in L_1, \forall i \in [n]$.

where $L_n, U_n \in \mathcal{V}_j$ and $L_1, U_1 \in \mathcal{V}_{j+1}$ for $j \in [m-1]$. Additionally if $(v_n, v_{k_1}), \ldots, (v_n, v_{k_q})$ are edges out of v_n, then add those corresponding edges between $L_n \in \mathcal{V}_j$ and $U_1 \in \mathcal{V}_{j+1}$ if $\mathcal{V}_j$ is a $\mathcal{V}_l$ gadget. Otherwise $\mathcal{V}_j$ is a $\mathcal{V}_u$ gadget and hence add the corresponding edges between $U_n \in \mathcal{V}_j$ and $L_1 \in \mathcal{V}_{j+1}$. It is easy to see that there is a bijection between walks of length m in G_A and paths of length m in $\mathcal{W}$. The gadget $\mathcal{W}$ that results is of constant treewidth.

5 Open Problems

What is the complexity of other linear algebraic invariants such as minimal polynomial of a bounded tree-width matrix? What is the complexity of counting Euler Tours in undirected tours in bounded treewidth graphs? On general graphs, this problem is known to be #P-complete[BW05]. See [CCM12,CCM13] for some recent progress on this problem.

References

AB09. Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

ABO99. Eric Allender, Robert Beals, and Mitsunori Ogihara. The complexity of matrix rank and feasible systems of linear equations. *Computational Complexity*, 8(2):99–126, 1999.

AEB87. T. Aardenne-Ehrenfest and N.G. Bruijn. Circuits and trees in oriented linear graphs. In Ira Gessel and Gian-Carlo Rota, editors, *Classic Papers in Combinatorics*, Modern Birkhuser Classics, pages 149–163. Birkhuser Boston, 1987.

Ber84. Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.*, 18(3):147–150, 1984.

BW05. Graham Brightwell and Peter Winkler. Counting eulerian circuits is# p-complete. In *ALENEX/ANALCO*, pages 259–262. Citeseer, 2005.

CCM12. Prasad Chebolu, Mary Cryan, and Russell Martin. Exact counting of euler tours for generalized series-parallel graphs. *J. Discrete Algorithms*, 10:110–122, 2012.

CCM13. Prasad Chebolu, Mary Cryan, and Russell Martin. Exact counting of euler tours for graphs of bounded treewidth. *CoRR*, abs/1310.0185, 2013.
CF12. Yijia Chen and Jorg Flum. On the ordered conjecture. In *Proceedings of the 2012 27th Annual IEEE/ACM Symposium on Logic in Computer Science*, pages 225–234. IEEE Computer Society, 2012.
CM87. Stephen A. Cook and Pierre McKenzie. Problems complete for deterministic logarithmic space. *J. Algorithms*, 8(3):385–394, 1987.
Cou90. Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and computation*, 85(1):12–75, 1990.
DDN13. Bireswar Das, Samir Datta, and Prajakta Nimbhorkar. Log-space algorithms for paths and matchings in k-trees. *Theory Comput. Syst.*, 53(4):669–689, 2013.
DK14. Samir Datta and Raghav Kulkarni. Space complexity of optimization problems in planar graphs. In *TAMC*, pages 300–311, 2014.
DKLM10. Samir Datta, Raghav Kulkarni, Nutan Limaye, and Meena Mahajan. Planarity, determinants, permanents, and (unique) matchings. *TOCT*, 1(3), 2010.
EJT10. Michael Elberfeld, Andreas Jakoby, and Till Tantau. Logspace versions of the theorems of bodlaender and courcelle. In *FOCS*, pages 143–152, 2010.
Ete97. Kousha Etessami. Counting quantifiers, successor relations, and logarithmic space. *J. Comput. Syst. Sci.*, 54(3):400–411, 1997.
HAB02. W. Hesse, E. Allender, and D.A.M. Barrington. Uniform constant-depth threshold circuits for division and iterated multiplication. *Journal of Computer and System Sciences*, 65:695–716, 2002.
MV97. Meena Mahajan and V. Vinay. Determinant: Combinatorics, algorithms, and complexity. *Chicago J. Theor. Comput. Sci.*, 1997, 1997.
Sta13. Richard P Stanley. *Algebraic Combinatorics*. Springer-Verlag New York, 2013.
Tod91. S. Toda. Counting problems computationally equivalent to the determinant. Technical Report CSIM 91-07, Dept of Comp Sc & Information Mathematics, Univ of Electro-Communications, Chofu-shi, Tokyo, 1991.
TS41. WT Tutte and CAB Smith. On unicursal paths in a network of degree 4. *The American Mathematical Monthly*, 48(4):233–237, 1941.
Val79. Leslie G Valiant. The complexity of computing the permanent. *Theoretical computer science*, 8(2):189–201, 1979.
vzGG13. Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra (3. ed.)*. Cambridge University Press, 2013.

Acknowledgement

We would like to thank Abhishek Bhrushundi, Arne Meier, Rohith Varma and Heribert Vollmer for illuminating discussions regarding this paper. Special thanks are due to Johannes Köbler and Sebastian Kuhnert who were involved in the initial discussions on the proof of Theorems 4, 6; to Stefan Mengel for proof reading the paper and finding a gap in a previous "proof" of Theorem 4; and to Raghav Kulkarni for suggesting proof strategies for Corollary 5 and Lemma 3; and to Sebastian Kuhnert for the proof of Proposition 2. Thanks are also due to anonymous referees for pointing out errors in a previous version of the paper, for greatly simplifying the proof of Corollary 3 and for pointing out the reference [CF12]. This research is partially funded by a grant from the Infosys Foundation.

A Background

A.1 Background on Complexity Classes

For all the standard complexity classes, we refer the reader to [AB09]. We define some non-standard complexity classes that we refer to here:

Let $f = (f_n)_{n \geq 0}$ where $f_n : \{0,1\}^n \to \mathbb{Z}$ be a family of integer valued functions. f is in the complexity class GapL if and only if there is a nondeterministic logspace machine M such that for every x, $f(x)$ equals the difference between the number of accepting and rejecting paths of M on input x.

$\mathsf{C{=}L}$ is the class of languages L for which there exists a GapL functions f that evaluates to zero on precisely those $x \in L$. The class $\mathsf{L}^{\mathsf{C{=}L}}$ consists of all languages decidable by a logspace machine at the base with oracle access to a $\mathsf{C{=}L}$-complete language. The AC^0 hierarchy over $\mathsf{C{=}L}$ is built as an AC^0 circuit with gates which can make queries to a $\mathsf{C{=}L}$ oracle and is known to collapse to $\mathsf{L}^{\mathsf{C{=}L}}$. For more information on the subtleties of implementing this hierarchy, see [ABO99].

B Split Graph and the Biadjacency matrix

We will often find it convenient to interpret an arbitrary matrix $A \in \mathbb{Q}_{n \times n}$ as representing both a weighted directed graph G_A and the weighted biadjacency matrix of an undirected bipartite graph H_A. We will exploit this duality between these two interpretations of a matrix to compute linear algebraic invariants.

Let $H_A = (R \cup C, E)$ be the bipartite graph on $2n$ vertices where A_{ij} represents an edge between the $i \in R$ to $j \in C$ of the appropriate weight. When we say H_A has bounded treewidth, we mean that the underlying undirected graph has bounded treewidth. If we consider A to represent a digraph G_A on n vertices, we can construct the *split* graph of G_A namely $\mathrm{Split}(G_A)$ – for each vertex v of G_A, we create v_{in} and v_{out}. If there is an edge between u and v in G_A, then we add an edge in $\mathrm{Split}(G_A)$ between u_{out} and v_{in}. For all vertices v, we add an edge (v_{in}, v_{out}). $\mathrm{Split}(G_A)$ is an undirected graph on $2n$ vertices. We have the following easy proposition:

Proposition 4. *G_A has bounded treewidth, iff Split(G_A) has bounded treewidth.*

Proof. We can construct the tree decomposition of $\mathrm{Split}(G_A)$ from that of G_A by introducing one new vertex for every vertex in the bag. This almost doubles the treewidth. The other directions follows from the fact that G_A can be obtained as minor of $\mathrm{Split}(G_A)$.

C Omitted Proofs from Section 3

C.1 Proof of Lemma 2

Proof. We first define the graph G' in which we add three vertices b_l, b_0, b_r for every bag B in the tree decomposition of G. We will use the following running

example, let A be a parent bag in the tree decomposition with left child B and right child C. The scheme we propose adds vertices a_l, a_0, a_r, b_l, b_0, b_r and c_l, c_0, c_r, with $a_l, a_0, a_r, b_r, c_l \in A$, $a_l, b_l, b_0, b_r \in B$ and $a_r, c_l, c_0, c_r \in C$.

We put an edge between every bag vertex $b \in V(G') \setminus V(G)$ and every vertex v sharing a bag with b. Clearly this is a valid tree decomposition of tree-width which is at most 6[5] greater than that of G.

We now traverse the tree T' using an Euler Traversal (see [CM87]). Every internal bag of the the tree is visited thrice in an Euler traversal - first when the traversal visits the bag the first time, second after exploring its left child, and finally third when it is done exploring its right child. For a bag B in the tree decomposition, if b_l, b_0, b_r are the vertices added, the vertices are visited in the following order in the Euler traversal of the tree: First b_l is visited, after which the left child of B is explored. Following this, the traversal returns to B via b_0, after which the right child of B is explored. Finally we visit B again via b_r and proceed to B's sibling via B's parent.

Define NXT_{Bag} to be the following ordering of the bag vertices: $a_l < b_l < b_0 < b_r < a_0 < c_l < c_0 < c_r < a_r$. In the case when b has children, they are explored between b_l and b_0 (left child) and b_0 and b_r (right child).

Next we extend the relation NXT_{Bag} to a relation NXT_1 on pairs which include vertices of G. For every vertex v occurring in bag A if v does not belong to the left child B of A but belongs to the right child C then add the tuple $a_0 < v$ to NXT_1 else we just add $a_l < v$. Symmetrically, we add $v < a_0$ or $v < a_r$ depending on whether v belongs to the left child B but not to the right child C, or not. If it does not belong to either child, we add the tuples $b_r < v$ and $v < c_l$.

Now consider the transitive closure NXT_1^* of NXT_1. This may well not be a total order on the vertices of G'. But we have the following:

Claim. Bag vertices are totally ordered under NXT_1^*. If two vertices of G are incomparable under NXT_1^*, then there must exist a common bag to which both must belong.

Proof. That bag vertices are totally ordered follows from the definition of Euler traversal of a tree. To see the other part of the claim let $u, v \in V(G)$ be unordered by NXT_1^*. Let the least common ancestor (LCA) bag of all the bags to which u belong be Q and the LCA bag of all bags containing v be R and further P is the LCA bag of Q, R. From the assumption that u, v do not belong to the same bag not all three P, Q, R can be the same, in particular Q, R are distinct. If P is distinct from Q, R then from NXT_1^* we know that: $q_r < p_0 < r_l$ (assuming, wlog that Q is a left descendant and R a right descendant of P). Also, $u < q_r$ (or possibly even $u < q_0 < q_r$) and $r_l < v$ (or even $r_l < r_0 < v$) are present in NXT_1. Thus $u < v$ is present in NXT_1^*. This leaves the case when $P = Q$

[5] Alternately, following scheme also works: Add all the 6 bag vertices(b_l, b_0, b_r and c_l, c_0, c_r) of the children to the parent bag A. Symmetrically, add the bag vertices a_l, a_0, a_r to both B and C. Note that this is a local operation and it increases the treewidth of every bag by atmost 9: 6 from the children and 3 from the parent. The final graph G' obtained this way has treewidth at most 12 more than G.

(the other case, $P = R$, is symmetric). Let P' be the bag containing u which is nearest along the tree to R. Without (much) loss of generality suppose R is a descendant of the left child of P'. Let R' be the left child of P' (R' may be the same as R). From the fact that R is the highest bag containing v, we know that $v < r_r$. Because u does not occur in either child of P' we have that $r'_r < u$. Now r_r is either the same as r'_r (if $R = R'$) or $r_r < r'_r$ (if R is a proper descendant of R'). In either case we get $v < u$ in the transitive closure of NXT_1^*. Other cases are similar.

We can find NXT_1^* in L by using [EJT10]. From this we find pairs $u, v \in V(G)$ which are incomparable under NXT_1^*. Next we order any unordered u, v in L by their binary values and add these tuples (i.e. one of $u < v$ and $v < u$ for every unordered u, v) to NXT_1 to yield NXT. Notice that the vertices of G are totally ordered under NXT^* and because of the claim above NXT is compatible with the tree decomposition T'

C.2 Proof of Lemma 3

Proof. We replace each edge (u, v) in G_A by a series-parallel graph with edge weights from $\{0, 1\}$ such that the sum of weights of all paths between u and v exactly equals the weight of (u, v) and the construction doesn't alter the sign of the cycle cover in which (u, v) is present.

We use a construction similar to [Val79] – If the weight of the edge (u, v) is an n-bit integer $w = w_n w_{n-1} \dots w_1$, we create a gadget of weight 2^i for each w_i and add edges (u, s_i) and (t_i, v). We add self loops on all the vertices in this gadget except u and v. We also ensure that all paths between u and v are of equal length (say l). If n' is the number of new vertices added, then a cycle cover of G_B will consist of the original cycles in G_A along with the $n' - l$ self loops that result from the series-parallel graph now representing (u, v). Since the sign of a monomial in the determinant is decided by the number of cycles (equivalently the number of heads in a cycle cover) corresponding to the permutation of the monomial, if we can ensure that $n' - l$ is always even, then the sign of the cycle cover will be $(-1)^k = (-1)^{n'-l+k} = (-1)^k$. This is easily taken care of by replacing v by a new vertex v', removing all the edges (t_i, v), introducing edges (t_i, v') and (v', v).

It remains to show that from the tree decomposition of G_A, we can obtain a tree decomposition of G_B in L. This is easily accomplished as follows: For an edge (u, v) in G_A, find the highest bag b in the tree where it occurs and add a new bag b_{uv} as a child to b which just contains u and v and the edge (u, v). Attach the tree decomposition of the series-parallel graph gadget corresponding to (u, v) in G_B as a child of b_{uv}. It is easy to verify that what results is still a tree decomposition – Every edge in B is present in some bag. If (u, v) is present in bags b_1 and b_2, it is present in all the bags in the unique path between b_1 and b_2 in the tree. Here one has to argue two cases: If (u, v) is an edge in the series-parallel gadget, then it is only present in the bags corresponding to the tree decomposition of the series parallel graph (children of b_{uv}). Else (u, v) must have been an edge in G_A. In this case, after appending the tree decomposition

of the series-parallel graph to b_{uv}, it still stays a tree decomposition. The size of any bag in this decomposition is $\max(k, 2)$ (since any series-parallel graph has treewidth at most 2).